

Targeted Learning Ensembles for Optimal Individualized Treatment Rules with Time-to-Event Outcomes

Iván Díaz^{*1}, Oleksandr Savenkov¹, and Karla Ballman¹

¹Division of Biostatistics, Weill Cornell Medicine.

February 16, 2017

Abstract

We consider estimation of an optimal individualized treatment rule (ITR) from observational and randomized studies when data for a high-dimensional baseline variable is available. Our optimality criterion is with respect to delaying time to occurrence of an event of interest (e.g., death or relapse of cancer). We leverage semiparametric efficiency theory to construct estimators with desirable properties such as double robustness, which means our estimators are optimal under consistent estimation of at least one of two high-dimensional nuisance parameters. We propose two estimators of the optimal ITR, which arise from considering two loss functions aimed at (i) directly estimating the conditional treatment effect (also known as the blip function), and (ii) recasting the problem as a weighted classification problem that uses the 0-1 loss function. Our estimators are *super learning* ensembles that minimize the cross-validated risk of a linear combination of estimators in a user-supplied library of candidate estimators. We prove oracle inequalities bounding the finite sample excess risk of the estimator. The bounds depend on the excess risk of the oracle selector and the bias in estimation of the nuisance parameters. These oracle inequalities imply asymptotic optimality of the estimated optimal ITR in the sense that one of the two following claims holds: the estimated optimal ITR is consistent, or it is equivalent with the oracle selector. In a randomized trial with uninformative censoring, we show that the value of the super learner based on (ii) achieves rate at least as fast as $\log n/n$, whereas the value of the super learner based on (i) achieves the slower rate $(\log n/n)^{1/2}$. We illustrate our methods in the analysis of a phase III randomized study testing the efficacy of a new therapy for the treatment of breast cancer.

1 Introduction

Individualized treatment rules play a fundamental role in the precision medicine model for healthcare, whereby medical decisions are targeted to the individual based on their expected

^{*}corresponding author: ild2005@med.cornell.edu

clinical response, instead of the traditional one-size-fits-all approach. Mathematically, ITRs are a function that maps an individual’s pre-treatment covariates into an optimal treatment choice. In this paper, we are concerned with learning the optimal ITR from data collected as part of an observational or randomized study, where optimality is defined as the maximum delay in the occurrence of an undesirable event (e.g., death or relapse).

Recent advances in biomedical imaging and gene expression technology produce large amounts of data that can be used to tailor treatment to very specific patient characteristics. Methods to estimate the optimal ITR when the dimension of the patient’s covariates is small and the optimal ITR is defined with respect to a single time-point outcome include the work of [Song and Pepe \(2004\)](#) and [Song and Zhou \(2011\)](#). In high-dimensional settings, the curse of the dimensionality and the counterfactual nature of the problem, by which only the outcome under the treatment actually received is observed, pose important challenges to estimate the optimal ITR. Recent developments include methods that optimize the prediction error of the outcome and methods based on weighted classification. In seminal work, [Qian and Murphy \(2011\)](#) propose a two step approach, in which outcome regression is estimated using a linear model with a large set of basis functions and ℓ_1 regularized least squares, and the optimal ITR is estimated as the treatment level that maximizes the predictions of the outcome regression. An important result of their work is a bound on the excess risk of the estimated ITR in terms of the excess risk of the ℓ_1 regularized estimator. Inspection of their theoretical results highlights the issue that estimation of the ITR through minimization of the outcome prediction risk may yield a suboptimal ITR. As a solution, [Zhao et al. \(2012\)](#) proposed to estimate the optimal ITR using an inverse probability weighted 0-1 loss function. They show that the expectation of the proposed loss function is equal to the mean potential outcome under the rule. Thus, solving the classification problem guarantees optimality of the estimated rule. Their methods focus on optimizing the risk of a support vector machine classifier ([Cortes and Vapnik, 1995](#)), using a convex surrogate loss function to alleviate the computational complexity associated to the non-convexity of the 0-1 loss. [Song et al. \(2015\)](#) use this classification set up to estimate a parametric decision function using a regularization term such as the ℓ_1 penalty. [Zhang et al. \(2012a,b, 2015\)](#) generalize the classification framework to allow the use of a doubly robust loss function, with a focus on low-dimensional parametric spaces for the nuisance parameters as well as the ITR. [Rubin et al. \(2012\)](#) also use doubly robust loss functions but improve the performance of the classification task by allowing for more flexible methods such as bagged decision trees and support vector machines. [Robins et al. \(2008a\)](#) propose a cross-validation scheme to select a data-adaptive estimator over a set of basis functions. They focus on a randomized trial, and show that their proposal achieves optimality under smoothness assumptions on the outcome regression. [McKeague and Qian \(2014\)](#) present an outcome regression approach that may be used in the context of functional covariate data.

Adaptations of the above methods to the problem of survival outcomes subject to informative censoring have also been proposed. In particular, [Zhao et al. \(2011\)](#); [Goldberg](#)

and Kosorok (2012) use Q-learning, relying on sequential support vector regressions, to estimate the optimal sequential treatment rule that optimizes a survival outcome under right-censoring. Geng et al. (2015) also tackle estimation of the optimal ITR in a survival setting, using ℓ_1 regularization for the outcome regression, under the assumption that censoring is independent of covariates and the outcome. Their decision functions are restricted to linear functions. Their method is somewhat restrictive since it assumes that censoring is independent of the outcome and the covariates. In addition, their methods are optimal for prediction but may yield suboptimal treatment rules (see Qian and Murphy, 2011). Zhao et al. (2015) generalize the weighted classification approach of Zhao et al. (2012) to allow for informative censoring and doubly robust loss functions, but their decision functions are restricted to support vector machines.

The individualized treatment rules discussed in this paper are single time-point versions of dynamic treatment rules (DTR, see e.g., Murphy et al., 2001; Murphy, 2003; Robins, 2004; van der Laan et al., 2005; Moodie et al., 2007; Robins et al., 2008b; Zhao et al., 2009; Moodie et al., 2012; Chakraborty et al., 2013; van der Laan and Luedtke, 2015). DTRs address the problem of dynamically allocating treatment in a longitudinal setting, where treatment decisions are made conditional on all available information, including previous treatment history. The literature in DTRs has grown rapidly in the last decades; Chakraborty and Murphy (2014) provide an excellent review of key developments in the field.

In this article, we propose two methods to construct an ensemble of decision functions for the optimal ITR. Our ensembles are linear combinations of estimators in a user-supplied library, where the coefficients in the linear combination are chosen to minimize the cross-validated risk. We propose to use a doubly robust loss function with roots in efficient estimation theory for marginal causal effects (Moore and van der Laan, 2009; Díaz et al., 2015). In our context, double robustness means that the estimated ITRs will have certain optimality properties under consistent estimation of at least one of two nuisance parameters: (a) the hazard of the outcome at each time point conditional on covariates and treatment, and (b) the hazard of censoring and the treatment mechanism.

The library of candidate estimators may contain any of the algorithms discussed in the previous paragraphs. In light of the “no free lunch” theorems of Wolpert (2002) for supervised learning, for any given dataset, our ensembles are expected to have better or equal generalization error than any of the individual candidates in the library. We provide a formal proof of this claim in the form of an oracle inequality, which bounds the excess risk of our estimator in terms of the excess risk of the oracle estimator, which is defined as the combination of estimators that would be chosen in a hypothetical world in which an infinite validation sample is available and at least one of the nuisance parameters is known. Our methods are developed under the assumption that censoring is at random (Rubin, 1987), which means that censoring is random within strata of treatment and baseline variables. We also assume that treatment is randomized within strata of the covariates, either by nature or by experimentation.

The finite sample bounds we present are inspired by developments in the targeted learning literature, which establish the optimality of cross-validation in estimator selection for high-dimensional parameters (van der Laan and Dudoit, 2003). Related to our work, Luedtke and van der Laan (2016) consider super learning ensembles for estimation of optimal DTRs in two time points, considering several loss functions. They present oracle inequalities for inverse probability weighted loss functions in randomized studies.

The paper is organized as follows. in Section 2 we present our motivating example and introduce the relevant notation. in Section 3 we describe the causal inference problem as well as the conditions under which identification is possible. Section 4 presents the main contributions of our paper: two proposed estimators for the optimal ITR along with their oracle inequalities, as well as their asymptotic consequences. Proofs are presented in the Supplementary Materials. In Section 5 we present the results of the data analysis in our illustrative example; we conclude with a few remarks and directions of future research in Section 6.

2 Motivating Example

Different types of human breast cancer tumors have been shown to have heterogeneous response to treatments (Perou et al., 2000; Sotiriou and Puztai, 2009). Amplification of ERBB2 gene and associated overexpression of human epidermal growth factor receptor (HER2) encoded by this gene occur in 25-30% of breast cancers (Slamon et al., 2001). HER2-positive breast cancer is an aggressive form of the disease and the prognosis for such patients is generally poor (Slamon, 1987; Seshadri et al., 1993). The clinical efficacy of adjuvant trastuzumab, a recombinant monoclonal antibody, in early stage HER2-positive patients was demonstrated by several large clinical trials (Perez et al., 2011; Romond et al., 2005). Despite significant improvement in disease-free and overall survival of patients treated with trastuzumab, about 20-25% patients relapse within 3-5 years (Perez et al., 2011). In this paper we use data from the North Central Cancer Treatment Group N9831 study, a phase III randomized clinical trial testing the addition of trastuzumab to chemotherapy in stage I-III HER2-positive breast cancer.

The total number of patients enrolled in the NCCTG N9831 trial was 3,505. Samples from 1,390 patients, for whom there was available tissue, were used to quantify mRNA from a custom codeset of 730 genes created by experts. The available baseline variables may be thus be categorized in three classes: demographic (e.g, race, age, ethnicity), clinical (e.g., tumor grade, tumors size, nodal status, hormone receptor status), and gene expression. Among the 1,390 patients, 483 received chemotherapy alone (control arm) and 907 patients received chemotherapy plus trastuzumab (treatment arm).

The clinical challenge is to identify genetic and demographic profiles for patients with HER2-positive breast cancer who are unlikely to benefit from adjuvant trastuzumab.

2.1 Data and Notation

Assume individuals are monitored at K equally spaced time points $t = \{1, \dots, K\}$ (months in our study). Let T denote a time-to-event outcome taking values in $\{1, \dots, K\} \cup \{\infty\}$, where $T = \infty$ represents no event occurring in the follow-up period. Let $C \in \{0, \dots, K\}$ denote the censoring time defined as the time at which the individual is last observed in the study, and let $C = K$, represent administrative censoring. Let $A \in \{0, 1\}$ denote study arm assignment, and let W denote a vector of baseline variables, which may include gene expression as well as demographic, comorbidity, and other clinical data. Denote $\mathbb{1}(X)$ the indicator variable taking value 1 if X is true and 0 otherwise. The observed data vector for each participant is $O = (W, A, \Delta, \tilde{T})$, where $\tilde{T} = \min(C, T)$, and $\Delta = \mathbb{1}\{T \leq C\}$ is the indicator that the participant's event time is observed (uncensored). For a random variable X , we let X take values on a set \mathbf{O} .

We assume the observed data vector for each participant i , denoted $O_i = (W_i, A_i, \Delta_i, \tilde{T}_i)$, is an independent, identically distributed draw from the unknown joint distribution P_0 on $(W, A, \Delta, \tilde{T})$. The empirical distribution of O_1, \dots, O_n is denoted with P_n . We assume $P_0 \in \mathcal{M}$, where \mathcal{M} is the nonparametric model defined as all continuous densities on O with respect to a dominating measure ν . We use P to denote a generic distribution $P \in \mathcal{M}$, and $E_0(X)$ to denote the expectation of X with respect to P_0 . For a function $f(o)$, we denote $Pf = \int f(o)dP(o)$, and $\|f\|^2 = P_0 f^2$. We use $a \lesssim b$ to denote that a is smaller or equal than b up to a universal constant.

We can equivalently encode a single participant's data vector O using the following longitudinal data structure:

$$O = (W, A, R_0, L_1, R_1, L_2, \dots, R_{K-1}, L_K), \quad (1)$$

where $R_t = \mathbb{1}\{\tilde{T} = t, \Delta = 0\}$ and $L_t = \mathbb{1}\{\tilde{T} = t, \Delta = 1\}$, for $t \in \{0, \dots, K\}$. For a random variable X , we denote its history through time t as $\bar{X}_t = (X_0, \dots, X_t)$. For a given scalar x , the expression $\bar{X}_t = x$ denotes element-wise equality. The corresponding vector (1) for participant i is denoted by $(W_i, A_i, R_{0,i}, L_{1,i}, R_{1,i}, L_{2,i}, \dots, R_{K-1,i}, L_{K,i})$.

Define the following indicator variables for each $t \geq 1$:

$$I_t = \mathbb{1}\{\bar{R}_{t-1} = 0, \bar{L}_{t-1} = 0\}, \quad J_t = \mathbb{1}\{\bar{R}_{t-1} = 0, \bar{L}_t = 0\}.$$

The variable I_t is the indicator based on the data through time $t - 1$ that a participant is at risk of the event being observed at time t ; in other words, $I_t = 1$ means that all the variables $R_0, L_1, R_1, L_2, \dots, L_{t-1}, R_{t-1}$ in the data vector (1) equal 0, which makes it possible that $L_t = 1$. Analogously, J_t is the indicator based on the outcome data through time t and censoring data before time t that a participant is at risk of censoring at time t . We define $J_0 = 1$.

Define the hazard function for survival at time $m \in \{1, \dots, K\}$:

$$h(m, a, w) = P_0(L_m = 1 | I_m = 1, A = a, W = w),$$

among the population at risk at time m within strata of study arm and baseline variables. Similarly, for the censoring variable C , define the censoring hazard at time $m \in \{0, \dots, K\}$:

$$g_R(m, a, w) = P_0(R_m = 1 | J_m = 1, A = a, W = w).$$

We use the notation $g_A(a, w) = P_0(A = a | W = w)$, $g = (g_A, g_R)$, and $\eta = (h, g_A, g_R)$. Let p_W denote the marginal distribution of the baseline variables W . We add the subscript 0 to p_W, g, h to denote the corresponding quantities under P_0 .

3 Treatment Effect, Identification, and Optimal Individualized Treatment Rules

Define the potential outcomes $T_a : a \in \{0, 1\}$ as the event times that would have been observed had study arm assignment $A = a$ and censoring time $C = K$ been externally set with probability one. For a restriction time $\tau \in \{1, \dots, K\}$ of interest, we define the restricted survival time under treatment arm $A = 1$ as $\min(T_a, \tau)$. For a transformation Z of W , the treatment effect within strata of the covariates Z may be defined in terms of the so-called full-data “blip” function (see e.g., [Robins, 1997](#)) of the restricted mean survival time (RMST):

$$\theta_c(z) = E\{\min(T_1, \tau) - \min(T_0, \tau) \mid Z = z\}.$$

The transformation Z may represent a subset of covariates (e.g., gene expression), or the whole vector W . We define the marginal treatment effect as $\theta_{c,m} = E\{\min(T_1, \tau) - \min(T_0, \tau)\}$.

The superscript c denotes a causal parameter, that is, a parameter of the distribution of the potential outcomes T_1 and T_0 . It can be shown (see [Díaz et al., 2015](#)) that $E\{\min(T_a, \tau) | Z = z\} = \sum_{t=0}^{\tau-1} S_c(t, a, z)$, where $S_c(t, a, z) = P(T_a > t \mid Z = z)$ is the survival probability corresponding to the potential outcome under assignment to arm $A = a$ within strata $Z = z$. As a result, $\theta_c(z)$ may be expressed as

$$\theta_c(z) = \sum_{t=1}^{\tau-1} \{S_c(t, 1, z) - S_c(t, 0, z)\}, \quad (2)$$

since $S_c(0, a, z) = 1$ for $a \in \{0, 1\}$ and for all z .

An individualized treatment rule d is a function that maps the covariate values z of a given participant to a personalized treatment decision in $\{0, 1\}$. The potential time to event under a rule d is defined as $T_d = d(z)T_1 + \{1 - d(z)\}T_0$. Accordingly, the RMST under a treatment rule that assigns treatment according to $d(z)$ is equal to

$$E\{\min(T_d, \tau)\} = E\{d(Z)[\min(T_1, \tau) - \min(T_0, \tau)]\} - E\{\min(T_0, \tau)\}.$$

Because the last term does not depend on $d(z)$, we define the *value* of the ITR d as

$$V_c(d) = E\{d(Z)[\min(T_1, \tau) - \min(T_0, \tau)]\} = E\{d(Z)\theta_c(Z)\}.$$

The above equation provides the basis for the construction of an optimal ITR as

$$d_c(z) = \arg \max_{d \in \mathcal{D}} V_c(d) = \mathbb{1}\{\theta_c(z) > 0\},$$

where $\mathcal{D} = \{d : \mathbf{Z} \rightarrow \{0, 1\}\}$ is the space of functions that map the range of Z into a treatment decision in $\{0, 1\}$. Note that we define optimality of an ITR with respect to the RMST, though other effect measures could also be used.

3.1 Identification of Parameters in Terms of Observed Data Generating Distribution P_0

In this section we show how the blip function $\theta_c(z)$, the value function $V_c(d)$, and the optimal ITR $d_c(z)$, which are defined above in terms of the distribution of potential outcomes, can be equivalently expressed as functions $\theta_0(z)$, $V_0(d)$, and $d_0(z)$ of the observed data distribution $P_0(W, A, \Delta, \tilde{T})$, under the assumptions (a)-(d) below. This is useful since the potential outcomes are not always observed, in contrast to the observed data vector $(W, A, \Delta, \tilde{T})$ for each participant, whose distribution we can make direct statistical inferences about.

Define the following assumptions:

- (a) $T = \mathbb{1}(A = 0)T_0 + \mathbb{1}(A = 1)T_1$ (*consistency*);
- (b) A is independent of T_a conditional on W , for each $a \in \{0, 1\}$ (*randomization*);
- (c) C is independent of T_a conditional on (A, W) , for each $a \in \{0, 1\}$ (*random censoring*);
- (d) $g_{A,0}(a, w) > 0$ and $g_{R,0}(t, a, w) < 1$ whenever the P_0 -density of W is positive at $W = w$, for each $a \in \{0, 1\}$ and $t \in \{0, \dots, \tau - 1\}$ (*positivity assumption*).

We make assumptions (a)-(d) throughout the manuscript. Assumption (a) connects the potential outcomes to the observed outcome. Assumption (b) holds by design in a randomized trial. Assumption (c), which is similar to that in [Rubin \(1987\)](#), means that censoring is random within strata of treatment and baseline variables (which we abbreviate as “random censoring”). Assumption (d) states that each treatment arm has a positive probability, and that every time point has a hazard of censoring smaller than one, within each baseline variable stratum $W = w$ with positive density under P_0 .

Denote the survival function for T at time $t \in \{1, \dots, \tau - 1\}$ conditioned on study arm a and baseline variables w by

$$S(t, a, w) = P(T > t | A = a, W = w). \quad (3)$$

Similarly, define the following function of the censoring distribution:

$$G(t, a, w) = P(C \geq t | A = a, W = w). \quad (4)$$

Under assumptions (a)-(d), we have $T \perp\!\!\!\perp C | A, W$ and therefore $S(t, a, w)$ and $G(t, a, w)$ have the following product formula representations:

$$S(t, a, w) = \prod_{m=1}^t \{1 - h(m, a, w)\}; \quad G(t, a, w) = \prod_{m=0}^{t-1} \{1 - g_R(m, a, w)\}. \quad (5)$$

The potential outcome survival function $S_c(t, a, z)$ can be equivalently represented in terms of the observed data distribution as

$$S(t, a, z) = E \left[\prod_{m=1}^t \{1 - h(m, a, W)\} \middle| Z = z \right], \quad (6)$$

for $t \in \{1, \dots, K\}$, $a \in \{0, 1\}$; equality of $S_c(t, a, z)$ and the above display follows from (5) and

$$\begin{aligned} S_c(t, a, z) &= P(T_a > t \mid Z = z) \\ &= E [P(T_a > t | W) \mid Z = z] \\ &= E [P(T > t | A = a, W) \mid Z = z] \\ &= E [S(t, a, W) \mid Z = z], \end{aligned}$$

where the third equality above follows from (a) and (b).

It follows from (2) that the causal parameter $\theta_c(z)$ is equal to the following observed-data blip function:

$$\theta(z) = \sum_{t=1}^{\tau-1} E \left\{ \prod_{m=1}^t \{1 - h(m, 1, W)\} - \prod_{m=1}^t \{1 - h(m, 0, W)\} \middle| Z = z \right\}. \quad (7)$$

Thus, the value $V_c(d)$ of a rule d is equal to $V(d) = E\{d(Z)\theta(Z)\}$, and the corresponding optimal treatment rule is equal to $d_0(z) = \mathbb{1}\{\theta_0(z) > 0\}$, where we denote the corresponding true quantities (i.e., quantities computed w.r.t. P_0) as $\theta_0(z)$, $V_0(d)$, and $d_0(z)$.

4 Estimation of the Blip Function and the Optimal Rule

In this section we discuss various estimators for $\theta_0(z)$, which can be mapped to estimators of the optimal ITR through $d_0(z) = \mathbb{1}\{\theta_0(z) > 0\}$. Our general strategy relies on the concept of *censoring unbiased transformation*, given in Definition 1 below. This concept was first introduced by Fan and Gijbels (1994) and is further discussed in Rubin and van der Laan (2007), among others.

Definition 1 (Unbiased transformation). $D : \mathcal{O} \rightarrow \mathbb{R}$ is referred to as an unbiased transformation for $\theta_0(z)$ if $E_0 \{D(O) \mid Z = z\} = \theta_0(z)$.

The above definition motivates the construction of estimators of $\theta_0(z)$ by regressing the transformation $D(O)$ on the covariates Z . A common complication in this step is that most unbiased transformations typically depend on unknown *nuisance* parameters which must be estimated prior to carrying out the analysis. For example, the inverse-probability transformation

$$D_g(O) = \sum_{t=1}^{\tau-1} \frac{(2A-1)\mathbb{1}\{\bar{R}_{t-1}=0, \bar{L}_t=0\}}{g_A(A, W)G(t, A, W)}$$

depends on the treatment and censoring mechanisms g_A and G . The consistency of an estimator of $\theta_0(z)$ thus depends on the ability to consistently estimate these nuisance parameters.

In this work, we focus on the *doubly robust* censoring unbiased transformation D_η defined in Lemma 1 below. In addition to being a doubly robust unbiased transformation for $\theta_0(z)$ (i.e., providing robustness to inconsistent estimation of one out of two nuisance parameters), D_η is an efficient estimating function in the non-parametric model in the sense that it may be used to construct efficient estimators of the marginal treatment effect $\theta_{c,m}$ (see e.g., [Díaz et al., 2015](#)).

Lemma 1 (Doubly robust censoring unbiased transformation). *Define*

$$D_\eta(O) = \sum_{m=1}^{\tau-1} [I_m Z(m, A, W) \{L_m - h(m, A, W)\} + S(m, 1, W) - S(m, 0, W)], \quad (8)$$

where $Z(m, A, W) = Z_1(m, A, W) - Z_0(m, A, W)$, and

$$Z_a(m, A, W) = - \sum_{t=m}^{\tau-1} \frac{\mathbb{1}\{A=a\}}{g_A(a, W)G(m, a, W)} \frac{S(t, a, W)}{S(m, a, W)}. \quad (9)$$

Assume $\eta = (h, g_A, g_R)$ is such that $h = h_0$ or $(g_A, g_R) = (g_{A,0}, g_{R,0})$. Then D_η is an unbiased transformation for $\theta_0(z)$, that is, $E_0\{D_\eta(O) \mid Z = z\} = \theta_0(z)$.

As a consequence of the previous lemma, the expected value of the quadratic loss function

$$L_\eta(O; \theta) = \{D_\eta(O) - \theta(Z)\}^2$$

is minimized at θ_0 if $\eta = (h, g)$ is such that either $h = h_0$ or $g = g_0$.

For a loss function L_η , we denote its expected value as $R_{0,\eta}(\theta) = E_0\{L_\eta(O; \theta)\}$ and refer to it as the risk. In the following section we discuss the construction of super learning ensembles of candidate estimators for θ that target minimization of the quadratic risk.

4.1 Super Learner Ensembles for the Blip Function

Consider a collection of estimation algorithms for estimating θ_0 , hereby called a *library*, $\mathcal{L} = \{\hat{\theta}_j : j = 1, \dots, J\}$. For an estimator $\hat{\eta}$ of η_0 , in light of the discussion of the previous section, this library may be constructed by considering any predictive algorithm that minimizes the quadratic risk for prediction of the doubly robust unbiased transformation $D_{\hat{\eta}}(O)$. The literature in machine and statistical learning provides us with a wealth of algorithms that may be used in this step. Examples include algorithms based on regression trees (e.g., random forests, Bayesian regression trees), algorithms based on smoothing (e.g., generalized additive models, local polynomial regression, multivariate adaptive regression splines), and others (e.g., support vector machines, neural networks).

Consider the following cross-validation set up. Let $\mathcal{V}_1, \dots, \mathcal{V}_K$ denote a random partition of the index set $\{1, \dots, n\}$ into K validation sets of approximately the same size. That is, $\mathcal{V}_k \subset \{1, \dots, n\}$; $\bigcup_{k=1}^K \mathcal{V}_k = \{1, \dots, n\}$; and $\mathcal{V}_k \cap \mathcal{V}_{k'} = \emptyset$. In addition, for each k , the associated training sample is given by $\mathcal{T}_k = \{1, \dots, n\} \setminus \mathcal{V}_k$. Denote $\hat{\eta}_k$ the estimator of η_0 trained only using data in \mathcal{T}_k . Likewise, denote by $\hat{\theta}_{j,k}$ the estimator of θ_0 obtained by training the j -th predictive algorithm in \mathcal{L} using only data in the sample \mathcal{T}_k (e.g., regressing $D_{\hat{\eta}_k}(O_i)$ on V_i for $i \in \mathcal{T}_k$). The cross-validated prediction risk of $\hat{\theta}_j$ is defined as

$$\hat{R}_{\hat{\eta}}(\hat{\theta}_j) = \frac{1}{K} \sum_{k=1}^K \frac{1}{|\mathcal{V}_k|} \sum_{i \in \mathcal{V}_k} L_{\hat{\eta}_k}(O_i; \hat{\theta}_{j,k}).$$

Alternatively, we sometimes use $k(i)$ to denote the index of the validation set that contains observation i , and rewrite

$$\hat{R}_{\hat{\eta}}(\hat{\theta}_j) = \frac{1}{K} \sum_{i=1}^n \frac{1}{|\mathcal{V}_{k(i)}|} L_{\hat{\eta}_{k(i)}}(O_i, \hat{\theta}_{j,k(i)}).$$

In this paper we consider an ensemble learner given by a convex combination

$$\hat{\theta}_\alpha(z) = \sum_{j=1}^J \alpha_j \hat{\theta}_j(z); \quad \alpha_j \geq 0; \quad \sum_{j=1}^J \alpha_j = 1.$$

The weights α_j are chosen to minimize the cross-validated risk of the above combination, that is:

$$\hat{\alpha} = \arg \min_{\alpha \in \mathcal{A}} \sum_{i=1}^n \frac{1}{|\mathcal{V}_{k(i)}|} \left\{ D_{\hat{\eta}_{k(i)}}(O_i) - \sum_{j=1}^J \alpha_j \hat{\theta}_{j,k(i)}(V_i) \right\}^2 \quad \text{subject to } \alpha_j \geq 0; \quad \sum_{j=1}^J \alpha_j = 1.$$

The above expression is a weighted ordinary least squares problem with constraints on the coefficients, and may therefore be solved using standard off-the-shelf regression or optimization software. We denote this super learner with $\hat{\theta}_{\text{sl}} = \hat{\theta}_{\hat{\alpha}}$.

The optimality of general cross-validation selection procedures is discussed in [van der Laan & S. Dudoit & A.W. van der Vaart \(2006\)](#); [van der Vaart et al. \(2006\)](#). Optimality here is defined in terms of asymptotic equivalence with the *oracle* risk, which we define as the risk computed when (i) one of the components of the nuisance parameter η_0 is known, and (ii) a validation sample of infinite size is available to assess the performance of the estimator. Specifically

Definition 2 (Oracle risk and oracle selector). *Let $\eta_1 = (g_1, h_1)$, where either $g_1 = g_0$, or $h_1 = h_0$. The oracle risk of a candidate $\hat{\theta}_\alpha$ is defined as*

$$\tilde{R}_{\eta_1}(\hat{\theta}_\alpha) = \frac{1}{K} \sum_{k=1}^K \int \left\{ D_{\eta_1}(o) - \hat{\theta}_{\alpha,k}(z) \right\}^2 dP_0(o).$$

The oracle selector is equal to

$$\tilde{\alpha} = \arg \min_{\alpha} \tilde{R}_{\eta_1}(\hat{\theta}_\alpha) \quad \text{subject to } \alpha_j \geq 0; \quad \sum_{j=1}^J \alpha_j = 1,$$

and the corresponding oracle blip function is denoted with $\hat{\theta}_{\text{or}} = \hat{\theta}_{\tilde{\alpha}}$.

Note that $\tilde{R}_{\eta_1}(\theta_0) = \int L_{\eta_1}(o; \theta_0) dP_0(o)$ is the optimal risk (with respect to the loss function L_{η_1} , which in light of Lemma 1 is a valid loss function) achieved by the true θ_0 . The following theorem provides a bound on the excess risk of the estimator $\hat{\theta}_{\text{sl}}$ and the excess risk of $\hat{\theta}_{\text{or}}$. The excess risk for a selector $\hat{\alpha}$ is defined as the difference between the oracle risk of the selector $\hat{\alpha}$ and the optimal risk, i.e.,

$$\mathcal{E}^2(\hat{\theta}_{\hat{\alpha}}) = E\{\tilde{R}_{\eta_1}(\hat{\theta}_{\hat{\alpha}}) - \tilde{R}_{\eta_1}(\theta_0)\} = E P_0(\hat{\theta}_{\hat{\alpha}} - \theta_0)^2,$$

where the expectation is taken over draws of O_1, \dots, O_n . We show that the above excess risk is bounded by two terms: one depending on the excess risk of the oracle selector $\mathcal{E}(\hat{\theta}_{\text{or}})$, and another one depending a doubly robust bias term associated to estimation of η_0 .

Theorem 1 (Oracle inequality for the super learner of the blip function). *Assume*

$$P_0\{g_{A,0}(A, W) > \epsilon\} = P_0\{g_{R,0}(t, A, W) < 1 - \epsilon\} = 1$$

for some $\epsilon > 0$ and for all $t \in \{0, \dots, \tau - 1\}$. Let $\eta_1 = (g_1, h_1)$ denote the limit of $\hat{\eta}$ as $n \rightarrow \infty$, and assume that either $g_1 = g_0$ or $h_1 = h_0$. Define

$$\begin{aligned} B_1(\hat{\eta}, \eta_0) &= E\|(\hat{g} - g_0)(\hat{h} - h_0)\| \\ B_2(\hat{\eta}, \eta_0) &= E\left\{ \mathbb{1}(g_1 = g_0)\|\hat{g} - g_0\| + \mathbb{1}(h_1 = h_0)\|\hat{h} - h_0\| \right\}. \end{aligned}$$

Then, for $\delta > 0$

$$\mathcal{E}(\hat{\theta}_{\text{sl}}) \leq (1 + 2\delta)^{1/2} \mathcal{E}(\hat{\theta}_{\text{or}}) + C_1 \{(1 + \log n)/n\}^{1/2} + C_2 B_1(\hat{\eta}, \eta_0) + C_3 (\log n/n)^{1/4} \{B_2(\hat{\eta}, \eta_0)\}^{1/2} \quad (10)$$

for constants C_1 , C_2 , and C_3 .

An immediate implication of Theorem 1 is the following corollary, which shows the asymptotic optimality of the super learner $\hat{\theta}_{\text{sl}}$ in the sense that one of the two following claims hold: (i) it is consistent, or (ii) it is asymptotically equivalent to the oracle selector.

Corollary 1 (Asymptotic optimality of the super learner). *Assume the conditions of Theorem 1. Assume also that $B_1(\hat{\eta}, \eta_0) \rightarrow 0$ and $B_2(\hat{\eta}, \eta_0) \rightarrow 0$ as $n \rightarrow \infty$. Then*

- (i) *If the library \mathcal{L} contains a consistent estimator of θ_0 such that $\lim_{n \rightarrow \infty} \mathcal{E}(\hat{\theta}_{\text{or}}) = 0$, then the super learner is asymptotically unbiased: $\lim_{n \rightarrow \infty} \mathcal{E}(\hat{\theta}_{\text{sl}}) = 0$.*
- (ii) *If, as is typical, none of the candidates in \mathcal{L} is consistent, such that $\lim_{n \rightarrow \infty} \mathcal{E}(\hat{\theta}_{\text{or}}) = \mathcal{E}_0 > 0$, then the super learner $\hat{\theta}_{\text{sl}}$ is asymptotically equivalent to the oracle selector in the sense that $\lim_{n \rightarrow \infty} \mathcal{E}(\hat{\theta}_{\text{sl}}) = \mathcal{E}_0$.*

The super learner $\hat{\theta}_{\text{sl}}$ may be used to construct an estimator of the optimal rule as $\hat{d}(z) = \mathbb{1}\{\hat{\theta}_{\text{sl}}(z) > 0\}$. The convergence rate of $\hat{\theta}_{\text{sl}}$ and $V(\hat{d})$ may be established using Theorem 1, provided the convergence rate of $B_j(\hat{\eta}, \eta_0) : j \in \{1, 2\}$ is known. The following remark discusses the rates in the particular case of a randomized study with no censoring.

Remark 1. *Assume that $P_0(\Delta = 1) = 1$ so that there is no censoring, and that $g_{A,0}(w) = q \in (0, 1)$, as would be the case in a randomized study. Then, a logistic regression fit of A on W containing at least an intercept would yield an estimator satisfying $\|\hat{g}_A - g_{A,0}\|^2 = O_P(n^{-1})$. Plugging in the true value $g_{R,0}(t, a, w) = 0$ for $\hat{g}_R(t, a, w)$, and assuming \hat{h} is inconsistently estimated yields $B_1(\hat{\eta}, \eta_0) = O(n^{-1/2})$ and $B_2(\hat{\eta}, \eta_0) = O(n^{-1/2})$. In this case, Theorem 1 teaches us that the excess risk of the super learner converges to that of the oracle selector at rate at least as fast as $(\log n/n)^{1/2}$, which is the parametric rate inflated by a factor of $(\log n)^{1/2}$. The latter factor represents a penalty due to model selection. Qian and Murphy (2011) prove that $V_0(\tilde{d}) - V_0(\hat{d}) \lesssim \{\mathcal{E}^2(\hat{\theta}_{\text{sl}}) - \mathcal{E}^2(\hat{\theta}_{\text{or}})\}^{1/2}$, where $\tilde{d}(z) = \mathbb{1}\{\hat{\theta}_{\text{or}}(z) > 0\}$ is the oracle rule. This yields the following convergence rate:*

$$V_0(\tilde{d}) - V_0(\hat{d}) = O_P\left((\log n/n)^{1/2}\right).$$

Though intuitive, optimizing the prediction risk for θ_0 , as done in this section, may not result in a rule $\hat{d}(z)$ that optimizes the value $V_0(d)$ (see an illustrating simulation in Luedtke and van der Laan, 2016). This may result in suboptimal treatment rules. To solve this problem, in the following section we present a method to construct a super learner

ensemble that directly targets optimization of $V_0(d)$. We provide an oracle inequality analogous to that given in Theorem 1. We show that the $(\log n/n)^{1/2}$ rate of Remark 1 can be improved to the almost parametric rate of $\log n/n$, under the additional assumption that all patient profiles considered in \mathbf{Z} are subject to effect modification in the sense that $\inf_{z \in \mathbf{Z}} |\theta_0(z)| > 0$.

4.2 Super Learner Ensembles for the Optimal Rule

In this section we introduce the concept of a decision function, which we define as a function $f : \mathbf{Z} \rightarrow \mathbb{R}$ that yields a treatment rule $d_f(z) = \mathbb{1}\{f(z) > 0\}$. In a slight abuse of notation we use $V(f)$ to refer to the value of the rule d_f . Note that any function f_0 such that $\text{sign}\{f_0(z)\theta_0(z)\} = 1$ has optimal value $V_0(d_0)$. This provides intuition on the benefits of directly optimizing the value of the loss function instead of the risk of the blip function: an inconsistent estimator of the blip function may provide an optimal rule, as long as its sign is correct. For a given rule d_f , in light of Lemma 1, we have that $V_0(f) = E_0\{d_f(Z)D_\eta(O)\}$ if η is such that either $h = h_0$, or $g = g_0$. Thus, a decision function that optimizes the value of the rule d_f may be found as

$$f_0 \in \arg \max_f \int d_f(z) D_\eta(o) dP_0(o).$$

For a binary value $b \in \{0, 1\}$ and any X we have (Zhang et al., 2012a)

$$bX = \mathbb{1}\{X > 0\}|X| - |X|\mathbb{1}[\mathbb{1}\{X > 0\} \neq b]. \quad (11)$$

Thus, the optimization problem may be recast as $f_0 \in \mathcal{F}_0$, where $\mathcal{F}_0 = \arg \min_f \int L_\eta(o; f) dP_0$ and

$$L_\eta(o; f) = |D_\eta(o)| \mathbb{1}[\mathbb{1}\{D_\eta(o) > 0\} \neq d_f(z)]. \quad (12)$$

Expression (12) is a weighted classification loss function in which we aim to classify the binary outcome $\mathbb{1}\{D_\eta(O) > 0\}$ based on data V , using the 0-1 loss function with weights given by $|D_\eta(O)|$. The objective is to classify an individual who benefits from treatment arm $A = 1$ (i.e., an individual with $D_\eta(O) > 0$) as requiring treatment (i.e., $d_f(Z) = 1$), while penalizing for the loss $|D_\eta(O)|$ incurred if the individual were misclassified.

In what follows we consider a library of algorithms for estimation of the decision function $\mathcal{L} = \{\hat{f}_j(z) : j, \dots, J\}$. In light of the discussion of the previous sections, the most natural choice for a decision function is the blip function $\hat{\theta}(z)$. However, we do not restrict our setup to functions with a blip interpretation. In addition to estimators of the blip function $\theta_0(z)$, the library may contain other decision functions such as the support vector machines proposed by Zhao et al. (2015) and the parametric decision functions of Bai et al. (2016).

We construct an ensemble of the decision functions as

$$\hat{f}_\alpha(z) = \sum_{j=1}^J \alpha_j \hat{f}_j(z); \quad \alpha_j \geq 0, \quad (13)$$

where, unlike $\hat{\theta}_\alpha(z)$, we do not restrict the coefficients to sum to one. We do this because different decision functions need not be in the same scale, unlike blip functions. In this way, we generate an ensemble optimal ITR as $\hat{d}_\alpha(z) = \mathbb{1}\{\hat{f}_\alpha(z) > 0\}$. As in the previous section, we define the super learner selector as

$$\hat{\alpha} = \arg \min_{\alpha \in \mathcal{A}} \sum_{i=1}^n \frac{1}{|\mathcal{V}_{k(i)}|} L_{\hat{\eta}_k(i)} \left(O_i, \hat{f}_{\alpha, k(i)} \right) \text{ subject to } \alpha_j \geq 0,$$

where $\hat{f}_{\alpha, k(i)}$ represents (13) with $\hat{f}_j(z)$ replaced by $\hat{f}_{j, k(i)}(z)$: the j -th decision function estimated using the training sample $\mathcal{T}_{k(i)}$. The super learner of the decision function is defined as $\hat{f}_{\text{sl}}(z) = \hat{f}_{\hat{\alpha}}(z)$, and the corresponding optimal ITR is defined as $\hat{d}_{\text{sl}} = \mathbb{1}\{\hat{f}_{\text{sl}}(z) > 0\}$.

For $\eta_1 = (g_1, h_1)$ such that either $g_1 = g_0$ or $h_1 = h_0$, the oracle risk of the decision function is defined as

$$\tilde{R}_{\eta_1}(\hat{f}) = \frac{1}{K} \sum_{k=1}^K \int L_{\eta_1}(o, \hat{f}_k) dP_0(o).$$

The oracle selector of α is thus defined as $\tilde{\alpha} = \arg \min_{\alpha} \tilde{R}_{\eta_1}(\hat{f}_\alpha)$, and we denote $\hat{f}_{\text{or}} = \hat{f}_{\tilde{\alpha}}$. The excess risk of an estimator \hat{f} is equal to

$$\mathcal{E}(\hat{f}) = E\{\tilde{R}_{\eta_1}(\hat{f}) - \tilde{R}_{\eta_1}(f_0)\} = V_0(f_0) - V_0(\hat{f}).$$

In Theorem 2 below, we provide bounds on $\mathcal{E}(\hat{f}_{\text{sl}})$ in terms of the excess risk of the oracle selector $\mathcal{E}(\hat{f}_{\text{or}})$ and the bias terms $B_1(\hat{\eta}, \eta_0)$ and $B_2(\hat{\eta}, \eta_0)$ defined in Theorem 1.

Theorem 2 (Oracle inequality for the super learner of the optimal ITR). *Assume the conditions of Theorem 1. In addition, assume $\inf_{z \in \mathbf{Z}} |\theta_0(z)| > 0$, and that $\hat{\alpha}$ is computed in a grid of size Mn^q for some $M > 0$, $q > 0$. Then, for $\delta > 0$*

$$0 \leq \mathcal{E}(\hat{f}_{\text{sl}}) \leq (1 + 2\delta) \mathcal{E}(\hat{f}_{\text{or}}) + C_1 \frac{1 + \log n}{n} + C_2 B_1(\hat{\eta}, \eta_0) + C_3 \sqrt{\frac{\log n}{n}} B_2(\hat{\eta}, \eta_0).$$

for constants C_1 , C_2 , and C_3 , where B_1 and B_2 are defined as in Theorem 1.

We have the following corollary establishing the asymptotic consequences of the above theorem.

Corollary 2 (Asymptotic optimality of the super learner). *Assume the conditions of Theorem 2, and assume that at least one of h_0 and g_0 is consistently estimated so that $B_1(\hat{\eta}, \eta_0) \rightarrow 0$ and $B_2(\hat{\eta}, \eta_0) \rightarrow 0$ as $n \rightarrow \infty$. Then*

- (i) *If the library \mathcal{L} contains an estimator that converges to a function f_0 such that $\text{sign}\{f_0(z)\theta_0(z)\} = 1$, then $\lim_{n \rightarrow \infty} \mathcal{E}(\hat{f}_{\text{or}}) = 0$, and the super learner $\hat{d}_{\text{sl}}(z) = \mathbb{1}\{\hat{f}_{\text{sl}}(z) > 0\}$ achieves the optimal value $V_0(d_0)$ asymptotically: $\lim_{n \rightarrow \infty} E\{V_0(\hat{d}_{\text{sl}})\} = V_0(d_0)$.*

(ii) If none of the candidates in \mathcal{L} is consistent, such that $\lim_{n \rightarrow \infty} \mathcal{E}(\hat{f}_{\text{or}}) = \mathcal{E}_{\phi,0} > 0$, then the super learner \hat{f}_{sl} is asymptotically equivalent to the oracle selector in the sense that $\lim_{n \rightarrow \infty} \mathcal{E}(\hat{f}_{\text{sl}}) = \mathcal{E}_{\phi,0}$.

The following remark provides the convergence rate of the optimal ITR in the special case in which the doubly robust bias terms $B_1(\hat{\eta}, \eta_0)$ and $B_2(\hat{\eta}, \eta_0)$ achieve parametric convergence rates (e.g., in a randomized trial with no censoring.)

Remark 2. Consider the set up of Remark 1. Theorem 2 teaches us that the value of the super learner \hat{f}_{sl} using the 0-1 loss function converges to that of the oracle selector at rate at least as fast as $\log n/n$, which is very close to the parametric rate of $n^{-1/2}$. That is

$$V_0(\hat{d}_{\text{or}}) - V_0(\hat{d}_{\text{sl}}) = O_P(\log n/n).$$

This is a faster rate than that achieved with the quadratic loss function considered in Section 4.1. This improvement comes at the price of the additional assumption that $\inf_{z \in \mathbf{Z}} |\theta_0(z)| > 0$. This assumption may be restrictive in settings in which some profiles z do not benefit from either treatment arm.

4.3 Using a Surrogate Loss Function for the 0-1 Loss

It is well known in the statistical learning literature that minimizing (12) is generally difficult due to the discontinuity and non-convexity of the 0-1 loss. A common approach to mitigate the issues arising from the discontinuity and non-convexity of the 0-1 loss function is to use surrogate loss functions, such as the logistic loss $\phi(x) = \log(1 + \exp(-x))$ or the hinge loss $\phi(x) = \max(1 - x, 0)$. We have the following result, which teaches us that any decision function $d_{f_0}(z)$ based on a decision function $f_0 \in \mathcal{F}_{\phi,0}$ has the same performance as the optimal ITR $d_0(z)$.

Lemma 2. Assume η is such that either $h = h_0$, or $g = g_0$. Define

$$\mathcal{F}_{\phi,0} = \arg \min_f \int L_{\phi,\eta}(o; f) dP_0(o), \quad (14)$$

where the surrogate loss $L_{\phi,\eta}$ is defined as

$$L_{\phi,\eta}(o; f) = |D_\eta(o)| \phi(f(z) [2\mathbb{1}\{D_\eta(o) > 0\} - 1]).$$

Then we have $\mathcal{F}_{\phi,0} \subseteq \mathcal{F}_0$, where $\mathcal{F}_0 = \arg \min_f \int L_\eta(o; f) dP_0$.

5 Estimating the Optimal Chemotherapy for Breast Cancer Patients in our Motivating Application

In this section we demonstrate the use of our methods in the study described in Section 2. In order to estimate and assess the performance of the estimated rule using different datasets, we split our data into training and validation datasets, of sizes 1000 and 390, respectively.

5.1 Estimators of h and g_R

According to our theoretical results, the optimality of the estimated treatment rules hinges upon consistent estimation of at least one of the nuisance parameters h and g_R . As a result, it is crucial to employ flexible methods capable of unveiling complex patterns which are not visible to the human eye. As demonstrated below in Section 5.3, simple parsimonious models such as the Cox proportional hazards or logistic regression fail to detect these complex relations in the data.

In order to accurately estimate the nuisance parameters, we use an ensemble learner known as the super learner for prediction (van der Laan et al., 2007). We train the ensemble separately using data from each treatment arm, in order to fully account for treatment-covariate interactions. Like our ITR ensembles, super learning predictors build a combination of candidate predictors that minimize a cross-validated user-supplied risk function. Since g_R and h are conditional probabilities, we focus on logistic regression ensembles and the negative log-likelihood loss function, using the R implementation in the SuperLearner package (Polley et al., 2016). The candidate estimators included in the ensembles are:

RF	Random Forests.
XGB	Extreme Gradient Boosting.
MLP	Multilayer Perceptron.
GLM	Logistic Regression.
MARS	Multivariate Adaptive Splines
LASSO	L_1 Regularized Logistic Regression

For RF, XGB, and MLP, the tuning parameters are tuned using data splitting with the aid of the R caret package (Kuhn et al., 2016). To avoid $p > n$, GLM and MARS are estimated with a variable screening algorithm which computes univariate t-statistics and keeps only the 50 variables with a larger value. The coefficients of each predictor in the ensemble, when trained in the complete dataset using 5-fold cross-validation, are presented in Table 1.

		RF	XGB	MLP	GLM	MARS	LASSO
\hat{g}_R	$A = 1$	0.100	0.056	0.000	0.000	0.312	0.532
	$A = 0$	0.000	0.301	0.000	0.000	0.294	0.405
\hat{h}	$A = 1$	0.023	0.050	0.000	0.000	0.237	0.691
	$A = 0$	0.154	0.086	0.098	0.060	0.123	0.480

Table 1: Coefficients of the super learner ensemble for estimation of g_R and h .

5.2 Candidate Estimators for the Optimal Treatment Rule

According to our discussion in Section 4, there are at least three types of estimators for the optimal rule d_0 . The first type is a simple substitution estimator, obtained through inspection of equation (7), which consists in regressing the random variable

$$\hat{B}(W) = \sum_{t=1}^{\tau-1} \{\hat{S}(t, 1, W) - \hat{S}(t, 0, W)\}$$

on V , where \hat{S} is the estimator of the survival function corresponding to the estimator \hat{h} described in Section 5.1. The second type is obtained through regression of the unbiased transformation $D_{\hat{\eta}}(O)$ on V . The third type of estimation methods is obtained based on equation (12), and is obtained by classifying the binary outcome $\mathbb{1}\{D_{\hat{\eta}}(O) > 0\}$ as a function of V , with weights given by $|D_{\hat{\eta}}(O)|$. Here, $\hat{\eta} = (\hat{g}_R, \hat{h})$, where the components of $\hat{\eta}$ are as described in Section 5.1. Any regression or supervised classification technique available in the statistical learning literature may be used as a candidate for solving these problems.

In our application, we focus on the following candidates for estimating d_0 :

B-Reg	Regression of $\hat{B}(W)$ on V using super learning with candidate learners as described in Section 5.1.
D-Reg	Regression of $D_{\hat{\eta}}(W)$ on V using super learning with candidate learners as described in Section 5.1.
D-Class-RF	Weighted random forest classification of $\mathbb{1}\{D_{\hat{\eta}}(O) > 0\}$ as a function of V .
D-Class-XGB	As above but using extreme gradient boosting.
D-Class-GLM	Weighted logistic regression for classification of $\mathbb{1}\{D_{\hat{\eta}}(O) > 0\}$ as a function of V .

According to our discussion in Sections 4.2 and 4.3, we also train four super learning ensembles of the above candidate estimators, using different loss functions:

SL-Reg	Regression ensemble minimizing the expected quadratic loss function $(D_{\eta}(O) - \theta(V))^2$.
SL-Class-01	Classification ensemble minimizing the expected 0-1 loss function.
SL-Class-Hinge	Classification ensemble with surrogate hinge loss function.
SL-Class-Log	Classification ensemble with surrogate log loss function.

The coefficients of each candidate estimator in each ensemble are presented in Table 2.

These coefficients were computed using the Subplex (Rowan, 1990) routine implemented in the NLOpt nonlinear-optimization R package. For improved robustness, the 0-1 loss was optimized using 1000 different random starting values.

	SL-Reg	SL-Class-0-1	SL-Class-Hinge	SL-Class-Log
D-Class-RF	0.000	0.005	0.000	0.000
D-Class-XGB	0.792	0.001	0.945	0.869
D-Class-GLM	0.000	0.031	0.000	0.000
D-Reg	0.007	0.017	0.001	0.006
B-Reg	0.201	0.947	0.054	0.125

Table 2: Coefficients of each candidate in each ensemble (standardized to sum one).

5.3 Assessing the Performance of The Estimated Treatment Rule

Once each rule is estimated using only data in the training dataset, its value $V(\hat{d})$ is assessed on the validation dataset. To that effect, we use the targeted minimum loss based estimator of the RMST proposed by Díaz et al. (2015) (See also Moore and van der Laan, 2011).

Figure 1 presents the estimated RMST obtained with each estimated rule, along with 95% confidence intervals. For comparison, we also present the value of two static rules of interest: never treat and always treat. As is clear from the figure, the best candidate learner in our application is B-Reg. All super learning ensembles yield a value that is similar to that of B-Reg, demonstrating the oracle property of the super learner. Treating patients according to the optimal rule yields a restricted mean survival of 157.1 (s.d. 3.1) months. In comparison with the “always treat” rule, which yields 151.2 (s.d., 3.3) months, the optimal rule improves mean patient survival by 6 months.

According to Table 2, only the super learning ensemble based on the 0-1 loss assigns a large weight to the optimal algorithm B-Reg. In fact, its RMST (see Figure 1) is identical to that of B-Reg. The other ensembles assign more weight to the second best algorithm, D-Class-XGB, and have slightly smaller RMST. This is in agreement with our theoretical findings that the best performance is obtained using the 0-1 loss function.

It is also worth noting that three of the estimated rules (D-Class-RF, D-Class-GLM, and D-Reg) yield a RMST smaller or equal than the RMST of the static rule “always treat”. In particular the D-Class-GLM rule, which is often advocated because it yields parsimonious rules (Zhang et al., 2015), yields suboptimal performance in our application.




Figure 1: RMST estimated in the validation set, for different estimated ITRs.

6 Concluding Remarks

In this article we develop a super learning ensemble estimator for the optimal individualized treatment rule when the objective is to delay occurrence of an outcome, such as death or relapse. We discuss several ways to generate candidate estimators, and present theoretical properties of our super learning ensemble, as well as an application to estimating the optimal treatment for breast cancer patients.

The convergence rate based on our proposed weighted classification approach is as fast as $\log n/n$, which is very close to the parametric $n^{-1/2}$ rate that would be obtained in a parametric model. This is remarkable since we make no parametric assumptions. This rate contrasts with the slower rate $(\log n/n)^{1/2}$ obtained when the ensemble is trained to minimize the proposed quadratic loss for regressing $D_\eta(O)$ on V . A limitation of our theoretical results is that the fast rate $\log n/n$ only applies when the classification procedure optimizes the 0-1 loss. Because the 0-1 loss is non-convex, the optimization problem is challenging. The question of whether the fast rate $\log n/n$ applies to the classification problem that uses surrogate convex loss functions is the subject of future research.

In comparison to Theorem 1, Theorem 2 has the additional assumption that the optimization of the loss function is carried out in a grid polynomial size in n . Inspection of the

proofs of the theorems in the appendix reveals the reason for the additional assumption: the 0-1 loss function is non-smooth and the Lipschitz condition used in the proof of Theorem 1 does not apply. As demonstrated in our data application, this assumption is likely to have little practical consequences, but it is unclear to us whether it can be removed.

In our application, we have decided to use data splitting to train and assess the performance of the estimated rules. Though correct, this approach may be unnecessary, since the value of the rule may be assessed using the training dataset, under certain conditions derived by [Luedtke et al. \(2016\)](#).

A Proofs of Theorems and Lemmas

A.1 Lemma 1

Proof For simplicity, consider the treatment-time-specific function

$$D_{m,a,\eta}(O) = - \sum_{t=1}^m \frac{\mathbb{1}\{A=a\}I_t}{g_A(a,W)G(t,a,W)} \frac{S(m,a,W)}{S(t,a,W)} \{L_t - h(t,a,W)\} + S(m,a,W),$$

and note that $D_\eta = \sum_{m=1}^{\tau-1} (D_{m,1,\eta} - D_{m,0,\eta})$. For a function $f(t,a,w)$ we denote $Pf(t) = \int f(t,a,w)dP(w)$. Conditioning first on W in the above display yields

$$E_0\{D_{m,a,\eta_0} \mid Z\} = E_0 \left\{ \prod_{t=1}^m \{1 - h_0(t)\} \mid Z \right\}.$$

Thus, we have

$$\begin{aligned}
& E_0(D_{m,a,\eta} \mid Z) - E_0 \left\{ \prod_{t=1}^m \{1 - h_0(t)\} \mid Z \right\} \\
&= E_0 \left[\sum_{t=1}^m -\frac{S(m)}{S(t)} \frac{g_{A,0}}{g_A} \frac{G_0(t)}{G(t)} S_0(t) \{h_0(t) - h(t)\} + \prod_{t=1}^m \{1 - h(t)\} - \prod_{t=1}^m \{1 - h_0(t)\} \mid Z \right] \\
&= \sum_{t=1}^m E_0 \left[-\frac{S(m)}{S(t)} \frac{g_{A,0}}{g_A} \frac{G_0(t)}{G(t)} S_0(t-1) \{h_0(t) - h(t)\} + S_0(t-1) \{h_0(t) - h(t)\} \frac{S(m)}{S(t)} \mid Z \right] \\
&= \sum_{t=1}^m E_0 \left[-\frac{S(m)}{S(t)} S_0(t-1) \{h_0(t) - h(t)\} \left\{ \frac{g_{A,0}}{g_A} \frac{G_0(t)}{G(t)} - 1 \right\} \mid Z \right] \\
&= \sum_{t=1}^m E_0 \left[-\frac{S(m)}{S(t)} S_0(t-1) \{h_0(t) - h(t)\} \left\{ \frac{g_{A,0}}{g_A G(t)} \{G_0(t) - G(t)\} + \frac{1}{g_A} (g_{A,0} - g_A) \right\} \mid Z \right] \\
&= \sum_{t=1}^m E_0 \left[-\frac{S(m)}{S(t)} S_0(t-1) \{h_0(t) - h(t)\} \left\{ \frac{g_{A,0}}{g_A G(t)} \sum_{k=0}^{t-1} G_0(k) \{g_{R,0}(k) - g_R(k)\} \frac{G(t)}{G(k+1)} \right. \right. \\
&\quad \left. \left. + \frac{1}{g_A} (g_{A,0} - g_A) \right\} \mid Z \right]
\end{aligned}$$

Plugging in $g = g_0$ or $h = h_0$ yields the result. \square

A.2 Theorem 1

Proof We start by assuming the minimization of the risk in the definition of $\hat{\alpha}$ and $\tilde{\alpha}$ is carried out in a grid $\mathcal{B}_n \subset \mathcal{B} = \{\alpha \in \mathbb{R}^J : \alpha_j \geq 0, \sum_{j=1}^J \alpha_j = 1\}$ of polynomial size in n (that is $|\mathcal{B}_n| \lesssim n^q$) for some $1 \leq q < \infty$, but do away with this assumption at the end of the proof. Let $\hat{\beta}$ and $\tilde{\beta}$ denote the cross-validated and oracle selectors when the risk minimization is performed in \mathcal{B}_n rather than \mathcal{B} . We use $P_{n,k}$ to denote the empirical distribution corresponding to the validation set \mathcal{V}_k , as well as $E_K(X) = K^{-1} \sum_{k=1}^K X_k$ to denote an average across validation splits. We denote $PL(\theta) = \int L(o; \theta) dP(o)$. Define the centered loss function

$$L_\eta^0(O; \theta) = L_\eta(O; \theta) - L_\eta(O; \theta_0),$$

and denote \hat{R}_η^0 and \tilde{R}_η^0 the corresponding cross-validated and oracle risks.

For notational convenience we denote $R(\beta) = R(\hat{\theta}_\beta)$. Let

$$\eta^* = \begin{cases} (g_0, \hat{h}_k) & \text{if } g_1 = g_0 \text{ and } h_1 \neq h_0 \\ (\hat{g}_k, h_0) & \text{if } g_1 \neq g_0 \text{ and } h_1 = h_0 \\ (g_0, h_0) & \text{if } g_1 = g_0 \text{ and } h_1 = h_0 \end{cases}$$

Note that $\tilde{R}_{\eta^*}^0(\beta) = \mathcal{E}(\hat{\theta}_\beta)$. For $\delta > 0$ we have

$$\begin{aligned}
0 &\leq \tilde{R}_{\eta^*}^0(\hat{\beta}) \\
&\leq \tilde{R}_{\eta^*}^0(\hat{\beta}) + (1 + \delta)\{\hat{R}_{\eta}^0(\tilde{\beta}) - \hat{R}_{\eta}^0(\hat{\beta})\} \\
&= (1 + 2\delta)\tilde{R}_{\eta^*}^0(\tilde{\beta}) \\
&\quad - (1 + \delta)\{\hat{R}_{\eta^*}^0(\hat{\beta}) - \tilde{R}_{\eta^*}^0(\hat{\beta})\} - \delta\tilde{R}_{\eta^*}^0(\hat{\beta}) \tag{15}
\end{aligned}$$

$$+ (1 + \delta)\{\hat{R}_{\eta^*}^0(\tilde{\beta}) - \tilde{R}_{\eta^*}^0(\tilde{\beta})\} - \delta\tilde{R}_{\eta^*}^0(\tilde{\beta}) \tag{16}$$

$$\begin{aligned}
&+ (1 + \delta)\{\hat{R}_{\eta}^0(\tilde{\beta}) - \hat{R}_{\eta^*}^0(\tilde{\beta})\} \\
&\quad - (1 + \delta)\{\hat{R}_{\eta}^0(\hat{\beta}) - \hat{R}_{\eta^*}^0(\hat{\beta})\} \tag{17}
\end{aligned}$$

where the second inequality is a consequence of the definition of $\hat{\beta}$ as the minimizer of $\hat{R}_{\eta}(\beta)$, and the last equality is the result of adding and subtracting some terms. Denote (15) with T , (16) with H , and (17) with $Q_{\tilde{\alpha}}$.

Note that the assumptions of the theorem imply that $P_0\{|D_{\eta^*}(O)| \leq M\} = 1$ for some constant M . This, together with Lemma 5 below, allow the application of Lemma 3 in [van der Laan and Dudoit \(2003\)](#) (see also pages 143-145 of [Dudoit and van der Laan, 2005](#)) to show that

$$E(T + H) \lesssim \frac{1 + \log n}{n}.$$

It remains to analyze $Q_{\hat{\beta}}$ and $Q_{\tilde{\beta}}$. First, we write $Q_{\beta} = Q_{1,\beta} + Q_{2,\beta}$, where

$$\begin{aligned}
Q_{1,\beta} &= (1 + \delta)E_K P_0(L_{\hat{\eta}_k}^0 - L_{\eta^*}^0)(\theta_{\beta,k}) \\
Q_{2,\beta} &= (1 + \delta)E_K(P_{n,k} - P_0)(L_{\hat{\eta}_k}^0 - L_{\eta^*}^0)(\theta_{\beta,k})
\end{aligned}$$

For $Q_{1,\beta}$, note that $R_{\eta^*}^0(\beta) = E_K P_0(\theta_0 - \theta_{\beta,k})^2$. This yields, for $\beta \in (\hat{\beta}, \tilde{\beta})$,

$$\begin{aligned}
E Q_{1,\beta} &= (1 + \delta)E E_K P_0(L_{\hat{\eta}_k}^0 - L_{\eta^*}^0)(\theta_{\beta,k}) \\
&= 2(1 + \delta)E E_K P_0(D_{\hat{\eta}_k} - D_{\eta^*})(\theta_0 - \theta_{\beta,k}) \\
&= 2(1 + \delta)\{E E_K P_0(D_{\hat{\eta}_k} - D_{\eta_0})(\theta_0 - \theta_{\beta,k}) - E E_K P_0(D_{\eta^*} - D_{\eta_0})(\theta_0 - \theta_{\beta,k})\}
\end{aligned}$$

Conditioning on W first, from the definition of η^* , Lemma 1 shows that the second term in the right hand side is zero. Conditioning on W first along with the proof of Lemma 1

and the Cauchy-Schwartz inequality also yields

$$\begin{aligned}
EQ_{1,\beta} &= 2(1+\delta)E E_K P_0(D_{\hat{\eta}_k} - D_{\eta_0})(\theta_0 - \theta_{\beta,k}) \\
&= \sum_{t=1}^m E E_K P_0(\theta_0 - \theta_{\beta,k}) \left[-\frac{S(m)}{S(t)} S_0(t-1)\{h_0(t) - h(t)\} \right. \\
&\quad \left. \left\{ \frac{g_{A,0}}{g_A G(t)} \sum_{k=0}^{t-1} G_0(k)\{g_{R,0}(k) - g_R(k)\} \frac{G(t)}{G(k+1)} + \frac{1}{g_A}(g_{A,0} - g_A) \right\} \middle| W \right] \\
&\leq 2(1+\delta) [E E_K P_0(\theta_0 - \theta_{\beta,k})^2]^{1/2} \|(\hat{g} - g_0)(\hat{h} - h_0)\| \\
&\lesssim \sqrt{E \tilde{R}_{\eta^*}^0(\hat{\beta})} B_1(\hat{\eta}, \eta_0)
\end{aligned}$$

where the last inequality follows from Lemma 5 in Appendix A.6 and the definition of $\tilde{\beta}$ as the minimizer of $\tilde{R}_{\eta^*}^0(\beta)$, and the second to last inequality follows from Cauchy-Schwartz applied to the norm defined by the inner product $\langle f_k, g_k \rangle = E E_K P_0 f_k g_k$. For $Q_{2,\beta}$, note that $(P_{n,k} - P_0)(L_{\hat{\eta}_k} - L_{\eta^*})(\theta_{\beta,k})$ is an empirical processes with index set \mathcal{B}_n , where the latter set is finite. We will apply the following inequality for empirical processes with finite index set:

$$E \max_{f \in \mathcal{F}} |(P_n - P_0)f| \lesssim \sqrt{\frac{\log |\mathcal{F}|}{n}} \|F\|, \quad (18)$$

where F is an envelope of \mathcal{F} . This result is a direct consequence of Lemma 19.38 of [van der Vaart \(1998\)](#). Note that the all functions in $\mathcal{F}_k = \{(L_{\hat{\eta}_k}^0 - L_{\eta^*}^0)(\theta_{\beta,k}) : \beta \in \mathcal{B}_n\}$ satisfy

$$\begin{aligned}
P_0(L_{\hat{\eta}_k}^0 - L_{\eta^*}^0)^2(\theta_{\beta,k}) &= P_0\{(D_{\hat{\eta}_k} - D_{\eta^*})^2(\theta_0 - \theta_{\beta,k})^2\} \\
&\lesssim P_0(D_{\hat{\eta}_k} - D_{\eta^*})^2 \\
&\lesssim B_2^2(\hat{\eta}, \eta_1),
\end{aligned}$$

where the second inequality follows from Lemma 5. Thus, the envelope F_k of \mathcal{F}_k is bounded by the same quantity. This, together with (18) shows

$$EQ_{2,\beta} \lesssim \sqrt{\frac{\log n}{n}} B_2(\hat{\eta}, \eta_1).$$

This proves

$$0 \leq E \tilde{R}_{\eta_0}^0(\hat{\beta}) \lesssim (1+2\delta) E \tilde{R}_{\eta_0}^0(\tilde{\beta}) + \frac{1+\log n}{n} + \sqrt{E \tilde{R}_{\eta_0}^0(\hat{\beta})} B_1(\hat{\eta}, \eta_1) + \sqrt{\frac{\log n}{n}} B_2(\hat{\eta}, \eta_1),$$

which is equivalent to $x^2 - bx \leq c$ for

$$\begin{aligned} x &= \sqrt{E\tilde{R}_{\eta_0}^0(\hat{\beta})} \\ b &= B_1(\hat{\eta}, \eta_1) \\ c &= (1 + 2\delta)E\tilde{R}_{\eta_0}^0(\tilde{\beta}) + \frac{1 + \log n}{n} + \sqrt{\frac{\log n}{n}}B_2(\hat{\eta}, \eta_1). \end{aligned}$$

The quadratic formula $x \leq (b + \sqrt{b^2 + 4c})/2$ implies $x \leq b + \sqrt{c}$, which yields

$$0 \leq \sqrt{E\tilde{R}_{\eta_0}^0(\hat{\beta})} \lesssim \sqrt{(1 + 2\delta)E\tilde{R}_{\eta_0}^0(\tilde{\beta})} + \sqrt{\frac{1 + \log n}{n}} + B_1(\hat{\eta}, \eta_0) + \left[\frac{\log n}{n}\right]^{1/4} \sqrt{B_2(\hat{\eta}, \eta_0)} \quad (19)$$

From our definitions and assumptions, the function $f(\beta) = R_{\eta_0}(\hat{\theta}_\beta)$ satisfies the Lipschitz condition

$$\|f(\beta) - f(\alpha)\|_\infty \lesssim \|\beta - \alpha\|_2,$$

where $\|\cdot\|_\infty$ denotes the supremum norm and $\|\cdot\|_2$ the Euclidean norm. Thus $f(\hat{\beta}) - f(\hat{\alpha})$ and $f(\tilde{\beta}) - f(\tilde{\alpha})$ are both bounded by n^{-q} , which allows us to replace $(\hat{\beta}, \tilde{\beta})$ by $(\hat{\alpha}, \tilde{\alpha})$ in (19), completing the proof of the theorem. \square

A.3 Corollary 1

Proof Take $\delta = n^{-\epsilon}$ for any $\epsilon > 0$. Claim (i) follows trivially by taking limits in the two sides of the inequality. Claim (ii) follows from the following argument. Taking limits in both sides yields $\lim_{n \rightarrow \infty} \mathcal{E}(\hat{\theta}_{\text{sl}}) \leq \lim_{n \rightarrow \infty} \mathcal{E}(\hat{\theta}_{\text{or}})$. By definition of $\tilde{\alpha}$, we have $\mathcal{E}(\hat{\theta}_{\text{or}}) \geq \mathcal{E}(\hat{\theta}_{\text{sl}})$ for all n , so that $\lim_{n \rightarrow \infty} \mathcal{E}(\hat{\theta}_{\text{sl}}) = \lim_{n \rightarrow \infty} \mathcal{E}(\hat{\theta}_{\text{or}}) = \mathcal{E}_0$. \square

A.4 Theorem 2

Proof The proof of this Theorem has the same steps as the proof of Theorem 1 and we will only provide a sketch. For convenience in the calculations we use the loss function

$$L_\eta(o; f) = -D_\eta(o)d_f(z) = -\mathbb{1}\{D_\eta(o) > 0\}D_\eta(o) + |D_\eta(o)|\mathbb{1}[\mathbb{1}\{D_\eta(o) > 0\} \neq d_f],$$

which is equivalent to the one used in the Theorem. Note that the conditions of the Theorem allow application of Lemma 4 below to obtain

$$E(T + H) \lesssim \frac{1 + \log n}{n}.$$

For $Q_{1,\alpha}$ and $Q_{2,\alpha}$ we get

$$\begin{aligned}
EQ_{1,\beta} &= (1 + \delta) E_K P_0(L_{\hat{\eta}_k}^0 - L_{\eta^*}^0)(f_{\alpha,k}) \\
&= 2(1 + \delta) E E_K P_0(D_{\hat{\eta}_k} - D_{\eta_0})(d_{f_{\alpha,k}} - d_0) \\
&= \sum_{t=1}^m E E_K P_0(d_{f_{\alpha,k}} - d_0) \left[-\frac{S(m)}{S(t)} S_0(t-1) \{h_0(t) - h(t)\} \right. \\
&\quad \left. \left\{ \frac{g_{A,0}}{g_A G(t)} \sum_{k=0}^{t-1} G_0(k) \{g_{R,0}(k) - g_R(k)\} \frac{G(t)}{G(k+1)} + \frac{1}{g_A} (g_{A,0} - g_A) \right\} \middle| W \right] \\
&\leq 2(1 + \delta) [E E_K P_0(\theta_0 - \theta_{\beta,k})^2]^{1/2} \|(\hat{g} - g_0)(\hat{h} - h_0)\| \\
&\lesssim B_1(\hat{\eta}, \eta_0).
\end{aligned}$$

For $Q_{2,\alpha}$, note that $(P_{n,k} - P_0)(L_{\hat{\eta}_k} - L_{\eta^*})(f_{\alpha,k})$ is an empirical processes with index set \mathcal{A}_n , where the latter set is the finite set with Mn^q points in which $\hat{\alpha}$ is computed. We will apply inequality (18). Note that the all functions in $\mathcal{F}_k = \{(L_{\hat{\eta}_k}^0 - L_{\eta^*}^0)(f_{\alpha,k}) : \alpha \in \mathcal{A}_n\}$ satisfy

$$\begin{aligned}
P_0(L_{\hat{\eta}_k}^0 - L_{\eta^*}^0)^2(\theta_{\alpha,k}) &= P_0\{(D_{\hat{\eta}_k} - D_{\eta^*})^2(d_{f_{\alpha,k}} - d_0)^2\} \\
&\lesssim P_0(D_{\hat{\eta}_k} - D_{\eta^*})^2 \\
&\lesssim B_2^2(\hat{\eta}, \eta_1),
\end{aligned}$$

where the second inequality follows from Lemma 5. Thus, the envelope F_k of \mathcal{F}_k is bounded by the same quantity. This, together with (18) shows

$$EQ_{2,\alpha} \lesssim \sqrt{\frac{\log n}{n}} B_2(\hat{\eta}, \eta_1).$$

This completes the proof. □

A.5 Lemma 2

Proof This is a direct application of Theorems 1 (part 3) and 2 of [Bartlett et al. \(2006\)](#). See also Theorem 5 of [Luedtke and van der Laan \(2016\)](#). □

A.6 Lemmas

Lemma 3. *Consider the assumptions of Theorem 1. Let $Z = L_{\eta_0}(O; \theta) - L_{\eta_0}(O; \theta_0)$. We have*

$$\text{Var}_0(Z) \lesssim E_0(Z)$$

Proof First, note that

$$Z = \{\theta_0(Z) - \theta(Z)\}\{2D_{\eta_0}(O) - \theta(Z) - \theta_0(Z)\}.$$

In light of Lemma 1 we have

$$E_0(Z) = E_0\{\theta_0(Z) - \theta(Z)\}^2.$$

Note that $P_0\{|2D_{\eta_0} - \theta(Z) - \theta_0(Z)| \leq 4 \max(M, C_1)\} = 1$. Thus

$$\begin{aligned} \text{Var}_0(Z) &\leq E_0(Z^2) \\ &= E\{\theta_0(Z) - \theta(Z)\}^2 \{2D_{\eta_0}(O) - \theta(Z) - \theta_0(Z)\}^2 \\ &\leq 16 \max(M^2, C_1^2) E_0(Z), \end{aligned}$$

which completes the proof of the lemma. \square

Lemma 4. *Consider the assumptions of Theorem 2. Let*

$$L_{\phi, \eta}(o, f) = D_{\eta}(o) d_f(z),$$

Let $Z = L_{\eta_1}(O; \theta) - L_{\eta_1}(O; \theta_0)$. We have

$$\text{Var}_0(Z) \lesssim E_0(Z)$$

Proof We have

$$\begin{aligned} E_0[Z^2] &= E_0|d_0(Z) - d_f(Z)|^2 D_{\eta_1}^2 \\ &\leq C E_0 \mathbb{1}\{d_0(Z) \neq d_f(Z)\} \\ &\leq C E_0 \frac{|\theta_0(Z)|}{\inf_z |\theta_0(Z)|} \mathbb{1}\{d_0(Z) \neq d_f(Z)\} \\ &\leq C_2 E_0 |\theta_0(Z)| \mathbb{1}\{d_0(Z) \neq d_f(Z)\} \\ &= C_2 E_0(Z). \end{aligned}$$

\square

Lemma 5. *For each $\hat{\eta} = (\hat{g}, \hat{h}) \rightarrow \eta_1 = (g_1, h_1)$ such that either $g_1 = g_0$ or $h_1 = h_0$ define*

$$\eta^* = \begin{cases} (g_0, \hat{h}) & \text{if } g_1 = g_0 \text{ and } h_1 \neq h_0 \\ (\hat{g}, h_0) & \text{if } g_1 \neq g_0 \text{ and } h_1 = h_0 \\ (g_0, h_0) & \text{if } g_1 = g_0 \text{ and } h_1 = h_0. \end{cases}$$

We have

$$P_0(D_{\hat{\eta}} - D_{\eta^*})^2 \lesssim B^2(\hat{\eta}, \eta_1),$$

with B^2 defined in Theorem 1.

Proof First let $g_1 = g_0$ and $h_1 \neq h_0$. Then $\eta^* = (g_0, \hat{h})$ and straightforward algebra shows

$$P_0(D_{\hat{\eta}} - D_{\eta^*})^2 \lesssim \|\hat{g} - g_1\|^2$$

Analogously, for $g_1 \neq g_0$ and $h_1 = h_0$ we have

$$P_0(D_{\hat{\eta}} - D_{\eta^*})^2 \lesssim \|\hat{h} - h_1\|^2.$$

Now, for $g_1 = g_0$ and $h_1 = h_0$ we get

$$P_0(D_{\hat{\eta}} - D_{\eta^*})^2 \lesssim \{\|\hat{h} - h_1\| + \|\hat{g} - g_1\|\}^2.$$

Putting these results together proves the lemma. \square

Lemma 6. For two sequences a_1, \dots, a_m and b_1, \dots, b_m we have

$$\prod_{t=1}^m (1 - a_t) - \prod_{t=1}^m (1 - b_t) = \sum_{t=1}^m \left\{ \prod_{k=1}^{t-1} (1 - a_k) (b_t - a_t) \prod_{k=t+1}^m (1 - b_k) \right\}.$$

Proof Replace $(b_t - a_t)$ by $(1 - a_t) - (1 - b_t)$ in the right hand side and expand the sum to notice it is a telescoping sum. \square

References

- Xiaofei Bai, Anastasios A Tsiatis, Wenbin Lu, and Rui Song. Optimal treatment regimes for survival endpoints using a locally-efficient doubly-robust estimator from a classification perspective. *Lifetime Data Analysis*, pages 1–20, 2016.
- Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- Bibhas Chakraborty and Susan A Murphy. Dynamic treatment regimes. *Annual review of statistics and its application*, 1:447, 2014.
- Bibhas Chakraborty, Eric B Laber, and Yingqi Zhao. Inference for optimal dynamic treatment regimes using an adaptive m-out-of-n bootstrap scheme. *Biometrics*, 69(3):714–723, 2013.
- C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20:273–297, December 1995.

- Iván Díaz, Elizabeth Colantuoni, and Michael Rosenblum. Improved efficiency in the analysis of randomized trials with survival outcomes. *arXiv preprint arXiv:1511.08404*, 2015.
- S. Dudoit and M.J. van der Laan. Asymptotics of cross-validated risk estimation in estimator selection and performance assessment. *Statistical Methodology*, 2(2):131–154, 2005.
- Jianqing Fan and Irène Gijbels. Censored regression: local linear approximations and their applications. *Journal of the American Statistical Association*, 89(426):560–570, 1994.
- Yuan Geng, Hao Helen Zhang, and Wenbin Lu. On optimal treatment regimes selection for mean survival time. *Statistics in medicine*, 34(7):1169–1184, 2015.
- Yair Goldberg and Michael R Kosorok. Q-learning with censored data. *Annals of statistics*, 40(1):529, 2012.
- Max Kuhn, Contributions from Jed Wing, Steve Weston, Andre Williams, Chris Keefer, Allan Engelhardt, Tony Cooper, Zachary Mayer, Brenton Kenkel, the R Core Team, Michael Benesty, Reynald Lescarbeau, Andrew Ziem, Luca Scrucca, Yuan Tang, Can Candan, and Tyler Hunt. *caret: Classification and Regression Training*, 2016. URL <https://CRAN.R-project.org/package=caret>. R package version 6.0-73.
- Alexander R Luedtke and Mark J van der Laan. Super-learning of an optimal dynamic treatment rule. *The international journal of biostatistics*, 12(1):305–332, 2016.
- Alexander R Luedtke, Mark J Van Der Laan, et al. Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *The Annals of Statistics*, 44(2):713–742, 2016.
- Ian W McKeague and Min Qian. Estimation of treatment policies based on functional predictors. *Statistica Sinica*, 24(3):1461, 2014.
- Erica EM Moodie, Thomas S Richardson, and David A Stephens. Demystifying optimal dynamic treatment regimes. *Biometrics*, 63(2):447–455, 2007.
- Erica EM Moodie, Bibhas Chakraborty, and Michael S Kramer. Q-learning for estimating optimal dynamic treatment rules from observational data. *Canadian Journal of Statistics*, 40(4):629–645, 2012.
- KellyL. Moore and MarkJ. van der Laan. Rcts with time-to-event outcomes. In *Targeted Learning*, Springer Series in Statistics, pages 259–269. Springer New York, 2011. ISBN 978-1-4419-9781-4.

- KL Moore and MJ van der Laan. Increasing power in randomized trials with right censored outcomes through covariate adjustment. *Journal of biopharmaceutical statistics*, 19(6): 1099–1131, 2009.
- S.A. Murphy, M.J. van der Laan, and J.M. Robins. Marginal mean models for dynamic treatment regimens. *Journal of the American Statistical Association*, 96:1410–1424, 2001.
- Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355, 2003.
- Edith A Perez, Edward H Romond, Vera J Suman, Jong-Hyeon Jeong, Nancy E Davidson, Charles E Geyer Jr, Silvana Martino, Eleftherios P Mamounas, Peter A Kaufman, and Norman Wolmark. Four-year follow-up of trastuzumab plus adjuvant chemotherapy for operable human epidermal growth factor receptor 2-positive breast cancer: Joint analysis of data from ncctg n9831 and nsabp b-31. *Journal of Clinical Oncology*, 29(25): 3366–3373, 2011.
- Charles M Perou, Therese Sørbye, Michael B Eisen, Matt van de Rijn, Stefanie S Jeffrey, Christian A Rees, Jonathan R Pollack, Douglas T Ross, Hilde Johnsen, Lars A Akslen, et al. Molecular portraits of human breast tumours. *Nature*, 406(6797):747–752, 2000.
- Eric Polley, Erin LeDell, and Mark van der Laan. *SuperLearner: Super Learner Prediction*, 2016. URL <https://CRAN.R-project.org/package=SuperLearner>. R package version 2.0-19.
- Min Qian and Susan A Murphy. Performance guarantees for individualized treatment rules. *Annals of statistics*, 39(2):1180, 2011.
- James Robins, Lingling Li, Eric Tchetgen, Aad van der Vaart, et al. Higher order influence functions and minimax estimation of nonlinear functionals. In *Probability and Statistics: Essays in Honor of David A. Freedman*, pages 335–421. Institute of Mathematical Statistics, 2008a.
- James Robins, Liliana Orellana, and Andrea Rotnitzky. Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in medicine*, 27(23):4678–4721, 2008b.
- James M Robins. Causal inference from complex longitudinal data. In *Latent variable modeling and applications to causality*, pages 69–117. Springer, 1997.
- James M Robins. Optimal structural nested models for optimal sequential decisions. In *Proceedings of the second seattle Symposium in Biostatistics*, pages 189–326. Springer, 2004.

- Edward H Romond, Edith A Perez, John Bryant, Vera J Suman, Charles E Geyer Jr, Nancy E Davidson, Elizabeth Tan-Chiu, Silvana Martino, Soonmyung Paik, Peter A Kaufman, et al. Trastuzumab plus adjuvant chemotherapy for operable her2-positive breast cancer. *New England Journal of Medicine*, 353(16):1673–1684, 2005.
- Thomas Harvey Rowan. *Functional Stability Analysis of Numerical Algorithms*. PhD thesis, Austin, TX, USA, 1990. UMI Order No. GAX90-31702.
- Daniel Rubin and Mark J van der Laan. A doubly robust censoring unbiased transformation. *The international journal of biostatistics*, 3(1), 2007.
- Daniel B Rubin, Mark J van der Laan, et al. Statistical issues and limitations in personalized medicine research with clinical trials. *The international journal of biostatistics*, 8(1):1–20, 2012.
- Donald B Rubin. *Multiple Imputation for Nonresponse in Surveys*. John Wiley & Sons, 1987.
- Ram Seshadri, FA Firgaira, DJ Horsfall, K McCaul, V Setlur, and P Kitchen. Clinical significance of her-2/neu oncogene amplification in primary breast cancer. the south australian breast cancer study group. *Journal of Clinical Oncology*, 11(10):1936–1942, 1993.
- Dennis J Slamon, Brian Leyland-Jones, Steven Shak, Hank Fuchs, Virginia Paton, Alex Bajamonde, Thomas Fleming, Wolfgang Eiermann, Janet Wolter, Mark Pegram, et al. Use of chemotherapy plus a monoclonal antibody against her2 for metastatic breast cancer that overexpresses her2. *New England Journal of Medicine*, 344(11):783–792, 2001.
- DJ Slamon. Human breast cancer: correlation of relapse and. *Science*, 3798106(177):235, 1987.
- Rui Song, Michael Kosorok, Donglin Zeng, Yingqi Zhao, Eric Laber, and Ming Yuan. On sparse representation for optimal individualized treatment selection with penalized outcome weighted learning. *Stat*, 4(1):59–68, 2015.
- Xiao Song and Margaret Sullivan Pepe. Evaluating markers for selecting a patient’s treatment. *Biometrics*, 60(4):874–883, 2004.
- Xiao Song and Xiao-Hua Zhou. Evaluating markers for treatment selection based on survival time. *Statistics in medicine*, 30(18):2251–2264, 2011.
- Christos Sotiriou and Lajos Pusztai. Gene-expression signatures in breast cancer. *New England Journal of Medicine*, 360(8):790–800, 2009.

- Mark J van der Laan and Alexander R Luedtke. Targeted learning of the mean outcome under an optimal dynamic treatment rule. *Journal of causal inference*, 3(1):61–95, 2015.
- M.J. van der Laan and S. Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. Technical report, Division of Biostatistics, University of California, Berkeley, November 2003.
- M.J. van der Laan, M.L. Petersen, and M.M. Joffe. History-adjusted marginal structural models & statically-optimal dynamic treatment regimens. *The International Journal of Biostatistics*, 1(1):10–20, 2005.
- M.J. van der Laan, E. Polley, and A. Hubbard. Super learner. *Statistical Applications in Genetics & Molecular Biology*, 6(25):Article 25, 2007.
- M.J. van der Laan & S. Dudoit & A.W. van der Vaart. The cross-validated adaptive epsilon-net estimator. *Statistics & Decisions*, 24(3):373–395, 2006.
- A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- A.W. van der Vaart, S. Dudoit, and M.J. van der Laan. Oracle inequalities for multi-fold cross-validation. *Statistics & Decisions*, 24(3):351–371, 2006.
- David H. Wolpert. *The Supervised Learning No-Free-Lunch Theorems*, pages 25–42. Springer London, London, 2002. ISBN 978-1-4471-0123-9.
- Baqun Zhang, Anastasios A Tsiatis, Marie Davidian, Min Zhang, and Eric Laber. Estimating optimal treatment regimes from a classification perspective. *Stat*, 1(1):103–114, 2012a.
- Baqun Zhang, Anastasios A Tsiatis, Eric B Laber, and Marie Davidian. A robust method for estimating optimal treatment regimes. *Biometrics*, 68(4):1010–1018, 2012b.
- Yichi Zhang, Eric B Laber, Anastasios Tsiatis, and Marie Davidian. Using decision lists to construct interpretable and parsimonious treatment regimes. *Biometrics*, 71(4):895–904, 2015.
- Ying-Qi Zhao, Donglin Zeng, Eric B Laber, Rui Song, Ming Yuan, and Michael Rene Kosorok. Doubly robust learning for estimating individualized treatment with censored data. *Biometrika*, 102(1):151–168, 2015.
- Yingqi Zhao, Donglin Zeng, A John Rush, and Michael R Kosorok. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118, 2012.

Yufan Zhao, Michael R Kosorok, and Donglin Zeng. Reinforcement learning design for cancer clinical trials. *Statistics in medicine*, 28(26):3294–3315, 2009.

Yufan Zhao, Donglin Zeng, Mark A Socinski, and Michael R Kosorok. Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer. *Biometrics*, 67(4): 1422–1433, 2011.